

Unbiased Sparse Subspace Clustering By Selective Pursuit

Hanno Ackermann and Michael Yang and Bodo Rosenhahn

Abstract

Sparse subspace clustering (SSC) is an elegant approach for unsupervised segmentation if the data points of each cluster are located in linear subspaces. This model applies, for instance, in motion segmentation if some restrictions on the camera model hold. SSC requires that problems based on the l_1 -norm are solved to infer which points belong to the same subspace. If these unknown subspaces are well-separated this algorithm is guaranteed to succeed.

The algorithm rests upon the assumption that points on the same subspace are well spread. The question what happens if this condition is violated has not yet been investigated. In this work, the effect of particular distributions on the same subspace will be analyzed. It will be shown that SSC fails to infer correct labels if points on the same subspace fall into more than one cluster.

1 Introduction

This paper considers unsupervised classification of data of points which reside on multiple unknown and low-dimensional subspaces. The problem is to decide which points belong to the same subspace. This *subspace clustering* arises in problems such as motion segmentation (Kanatani 2001; Vidal, Ma, and Sastry 2005; Elhamifar and Vidal 2013), hand written digit clustering (Zhang et al. 2012), face clustering (Ho et al. 2003) and compression (Hong et al. 2005).

Sparse subspace clustering (SSC) (Elhamifar and Vidal 2013) estimates the self-expressiveness within the data: which points can be used to linearly approximate a point in question? This is performed by minimizing l_1 -norm problems. The support of each point is used to define the edge weights of a graph whose vertices correspond to the data points. Graph segmentation then reveals the class membership. In (Soltanolkotabi and Candes 2012), SSC was theoretically analyzed and bounds for successful segmentation were derived.

Similar works use the nuclear norm instead of the l_1 -norm to infer edge weights (Liu, Lin, and Yu 2010). An extension of SSC jointly estimates the parameters of a global subspace which includes all the data (Patel, Nguyen, and Vidal 2013). In a recent work, both steps of sparse optimization and spectral clustering were combined into a single, iterative algorithm (Li and Vidal 2015).

A problem stems from the sparsity of the affinity matrix

of the graph. As the edge weights are defined by the sparse coefficients, only few weights are known. The question then is whether the vertices of points on the same subspace are always connected. It was first raised in (Nasihatkon and Hartley 2011). There, it was concluded that connectivity is guaranteed in subspaces of dimension 2 and 3, yet not for higher dimensions. For subspaces of dimension 4 an example was given in which the graph was disconnected.

In this paper we focus on the same question. Different prior works, we analyze the case that points on the same subspace are not evenly distributed as required in (Soltanolkotabi and Candes 2012), but fall into different clusters. It will be shown that a few edges indeed connect vertices of different clusters. However, it will also be shown that the weights of such edges are negligible small compared with edges connecting vertices of the same cluster. In other words, the *relative* connectivity is too low to be significant. The reason is a bias of l_1 -norm based estimators.

The graph constructed from the sparse coefficients then consists of more disconnected subgraphs than the correct yet unknown number of subspaces. Since the subspace model is often motivated by some underlying physical model, for instance each rigidly moving body induces a 4-dimensional subspace in motion segmentation, over-segmenting the data can be difficult to correct. On the other hand, because most algorithms of this class (Elhamifar and Vidal 2013; Patel, Nguyen, and Vidal 2013; Li and Vidal 2015; Liu, Lin, and Yu 2010) require the number of clusters to be known in advance, cluster centers will be incorrect which causes the labelling to be wrong. It is therefore more desirable to address the original problem in the first place.

This paper is structured as follows: In Sec. 2, we summarize the notation used in this paper. In Sec. 3, the original sparse subspace clustering algorithm is shortly explained. An example of the effect of non-uniformly distributed points is given in Sec. 4. A formal analysis on the magnitude of the sparse coefficients is given in Sec. 5. The proposed algorithm is introduced in Sec. 6. Experimental results are shown in Sec. 7. The paper concludes with a summary in Sec. 8.

2 Notation

Bold lower-case letters \mathbf{x} indicate vectors, normal capital letters matrices, e.g. $X = [\mathbf{x}_1 \cdots \mathbf{x}_N]$, or constants, e.g. N . By $\mathbf{x}(j)$, we mean the j th element of vector \mathbf{x} ,




Figure 1: Upper row: two sets of points (red and blue) separated by angles of 2° , 4° , 10° and 20° from left to right. Lower row: corresponding connectivity matrices $|C| + |C^\top|$ resulting from solving Eq. (2) with $\lambda = 0.01$. The columns/rows correspond with the numbering of the points above.

The larger the angle between the two sets of points, the lower the connectivity between the two clusters as the magnitude of the entries on the anti-diagonal becomes smaller and smaller. For an angle of 20° , the two sets almost form two disconnected clusters.

and by $A(i, j)$ the (i, j) th entry of matrix A . Sparse subspace clustering rests upon the assumption that there are L subspaces $\mathcal{S}^{(l)} \in \mathbb{R}^m$, $l = 1, \dots, L$ with N_l points $X^{(l)} = [\mathbf{x}_1^{(l)}, \dots, \mathbf{x}_{N_l}^{(l)}]$ on it. Let $d^{(l)}$ denote the dimension of $\mathcal{S}^{(l)}$. Let $X = [X^{(1)} \dots X^{(L)}]$ denote the matrix of all $N_1 + \dots + N_L = N$ data points. Assume without loss of generality that X is ordered. The points \mathbf{x}_j are supposed to have unit length $\|\mathbf{x}_j\|_2 = 1$.

Let $T^{(l)}$ indicate the set of indices of all the points $\mathbf{x}^{(l)} \in \mathcal{S}^{(l)}$. Let $S_j \subseteq T^{(l)}$ indicate the support set of point \mathbf{x}_j , $j = \{1, \dots, N\}$. The cardinality of a set is denoted by $|\cdot|$. The k th point of the support set S_j of $\mathbf{x}_j^{(l)}$ is denoted by $\mathbf{x}_{j,k}^{(l)}$. Let \mathbf{y} be any of the points of X . The matrix $X_{-\mathbf{y}}$ indicates the matrix X without the column corresponding to \mathbf{y} .

By $\|\cdot\|_{\{0,1,2,F\}}$ we indicate the l_0 -pseudo-norm, or the l_1 -, l_2 -, or Frobenius-norm of the argument.

3 Sparse Subspace Clustering

Sparse subspace clustering (SSC) (Elhamifar and Vidal 2013) is based on the fact that any point $\mathbf{x}_j \in \mathcal{S}^{(l)}$ can be expressed by a linear combination of $d^{(l)}$ other points in the same subspace. Since such linear combinations do not use points of other classes, they can be used to infer the unknown class labels.

The problem is to compute the least number of points \mathbf{x}_k , $k \neq j$ such that a linear combination of the \mathbf{x}_k yields \mathbf{x}_j . Denote by \mathbf{c}_j the vector of mixing coefficients with its j th entry being equal to zero such that $\mathbf{x}_j = X\mathbf{c}_j$ holds true. The task to estimate the vector \mathbf{c}_j can be solved by minimizing

$$\min \|\mathbf{c}_j\|_0 \quad \text{s.t.} \quad \|X\mathbf{c}_j - \mathbf{x}_j\|_2^2 = 0 \quad \text{and} \quad c_{jj} = 0. \quad (1)$$

Since optimizing the l_0 -norm is difficult, it is often approximated by using the l_1 -norm. Allowing for data points \mathbf{x}_j contaminated by noise, it is possible to instead optimize

$$\min \|\mathbf{c}_j\|_1 \quad \text{s.t.} \quad \|X\mathbf{c}_j - \mathbf{x}_j\|_2^2 \leq \lambda \quad \text{and} \quad c_{jj} = 0. \quad (2)$$

for a scalar $\lambda > 0$.

Given the matrix $C = [\mathbf{c}_1 \dots \mathbf{c}_N]$, class labels can be inferred by means of spectral clustering of the graph $G = (V, E)$ with $|V| = N$ vertices corresponding to the points \mathbf{x}_j and edge weights E_{ij} defined by the matrix $|C| + |C^\top|$. These steps are the basis of the so-called sparse subspace clustering (SSC) algorithm proposed in (Elhamifar and Vidal 2013).

4 Subspace Connectivity in \mathbb{R}^2

Since only few edge weights are known, the authors of (Nashatkon and Hartley 2011) raised the question whether connectivity between vertices of the same cluster is always guaranteed. In other words, is the subgraph consisting of the vertices of one particular cluster connected? Their answer was that connectivity is guaranteed – and thus a correct result of SSC – if the subspace dimension is $d = 2$ or $d = 3$, but an example was given for $d = 4$ where connectivity is violated. Independently, the authors of (Soltanolkotabi and Candes 2012) concluded that SSC clusters correctly as long as the points are well spread within each cluster.

In this work, we will analyze the connectivity if the data is not well spread across each subspace but forms more or less well isolated clusters in a particular subspace. As an example, consider Fig. 1. Here, the upper row shows plots of two sets of points (depicted in red and blue). From left to right, the angle between the two sets increases from 2° to 4° and 10° until 20° . The bottom row of Fig. 1 shows the corresponding affinity matrices $|C| + |C^\top|$ resulting from optimizing Eq. (2) with $\lambda = 0.01$. The columns and rows correspond to the numbering of the points in the upper row of Fig. 1. It can be seen that the $(3, 4)$ and $(4, 3)$ entries have the same magnitudes as the $(2, 3)$, $(3, 2)$ and $(4, 5)$, $(5, 4)$ entries for an angle of 2° . However, at an angle of 20° , the entries on the anti-diagonal have almost negligible magnitude as compared to the other entries on the diagonal. Therefore, the

resulting graph consists of two almost disconnected clusters.

Apparently, the question raised in (Nasihatkon and Hartley 2011), namely whether there is connectivity, i.e. are there non-zero entries in the affinity matrix, is not sufficient. The affinities shown on the anti-diagonal of the rightmost affinity are > 0 yet their magnitude is negligible as compared to the other entries.

This example shows that spurious clusters can emerge if sparse coefficients are used as affinities. The example is equivalent to points on an affine line ($d = 2$) normalized to unit length. It therefore readily generalizes to higher-dimensional structures in \mathbb{R}^m , $m > 2$. The remainder of this paper focuses on the question how gaps between data on the same subspace influence the solutions \mathbf{c}_j of Eq. (2) and thus the *relative* connectivity. The relative connectivities then determine if the corresponding vertices of the graph form more or less disconnected clusters.

5 The Bias of l_1 -Norm Estimators

The analysis in this section solely concentrates on points on a single subspace $S^{(l)}$. For the sake of simplicity, the superscript (l) is omitted in the following.

Let the angle between two points \mathbf{x}_{j_1} and \mathbf{x}_{j_2} , $j_1 \neq j_2$ be defined by the inverse cosine of the scalar product of two unit length vectors, $\angle(\mathbf{x}_{j_1}, \mathbf{x}_{j_2}) = \arccos(\mathbf{x}_{j_1}^\top \cdot \mathbf{x}_{j_2})$. Assume without loss of generality that the points $\mathbf{x}_{j,k}$ are ordered such that $\angle(\mathbf{x}_j, \mathbf{x}_{j,1}) \geq \dots \geq \angle(\mathbf{x}_j, \mathbf{x}_{j,K})$.

Noiseless l_1 -Norm: $\lambda = 0$

Given a point \mathbf{y} , a matrix $X_{-\mathbf{y}}$ and the K points $\mathbf{x}_{y,k}$, $k = 1, \dots, K$ of a support set S_y of \mathbf{y} , and let \mathbf{c}_y be a solution to $\|X_{-\mathbf{y}}\mathbf{c}_y - \mathbf{y}\|_2^2 \leq \lambda$ so that all entries but those corresponding to the points $\mathbf{x}_{y,k}$ in $X_{-\mathbf{y}}$ are zero. The question then is how does $\|\mathbf{c}_y\|_1$ change if a different support set $S'_y \subseteq T_y$ is selected.

The geometry of this case is shown in the left plot of Fig. 2. The intersection of the red dash-dotted line with the line through the point $\mathbf{x}_{y,1}$ indicates the point $\mathbf{c}_y(1)\mathbf{x}_{y,1}$. The point \mathbf{s} indicates the intersection of this line with the subspace $\text{span}(Q)$ spanned by the remaining points $Q = [\mathbf{x}_{y,2} \dots]$. It can be seen that $\|\mathbf{s}\|_1$ grows as the angle θ between $\text{span}(Q)$ and $\text{span}(\mathbf{x} - \mathbf{y})$ decreases. This intuition motivates the following proposition:

Proposition 1. *The l_1 -norm of \mathbf{c}_y increases with the angle $\angle(\mathbf{y}, \mathbf{x}_{y,1})$ between \mathbf{y} and $\mathbf{x}_{y,1}$*

$$\|\mathbf{c}_j\|_1 \geq \cot\left(f(\angle(\mathbf{y}, \mathbf{x}_{y,1})) + \rho - \frac{\pi}{2}\right) \cdot \eta. \quad (3)$$

Proof. Let $\mathbf{v}_1 = (c_{j,1}\mathbf{x}_{y,1} - \mathbf{y})/\|c_{j,1}\mathbf{x}_{y,1} - \mathbf{y}\|_2$, $\mathbf{v}_2 = (P_Q\mathbf{y} - \mathbf{y})/\|P_Q\mathbf{y} - \mathbf{y}\|_2$ where P_Q is the orthogonal projector onto $\text{span}(Q)$, and $f(\angle(\mathbf{y}, \mathbf{x}_{y,1})) = \cos^{-1} \mathbf{v}_1^\top \mathbf{v}_2$. Let $\alpha = \cos^{-1} \mathbf{v}_1^\top \mathbf{y}$ be the angle between the line connecting \mathbf{y} and $c_{j,1}\mathbf{x}_{y,1}$. Assuming that $|T_y| > |S_y|$, the support set $\mathbf{x}_{y,2}, \dots$ is selected such that $\text{span}(Q) \perp \text{span}([\mathbf{y} \dots \mathbf{x}_{y,1}])$ in order to minimize $\|\mathbf{x}\|_2$. Then, we can define the angles $\alpha_1 =$

$\cos^{-1} -\mathbf{y}^\top \mathbf{v}_2$, $\alpha_2 = \pi - \alpha - \alpha_1$, and $\theta = \alpha + \alpha_1 - \pi/2$ which motivates

$$\|\mathbf{x} - P_Q\mathbf{y}\|_2 = \cotan \theta \cdot \|(I - P_Q)\mathbf{y}\|_2. \quad (4)$$

If we now define $\boldsymbol{\nu} = (I - P_Q)\mathbf{y}$, $\eta = \|\boldsymbol{\nu}\|_2$, and $\rho = \cos^{-1}(\mathbf{y}^\top \boldsymbol{\nu})/\eta$ we arrive at the claim since $\|\mathbf{y}\|_1 \geq \|\mathbf{y}\|_2$. \square

Given \mathbf{y} , and d points $\mathbf{x}_{y,1}, \dots, \mathbf{x}_{y,d}$ of a support set S_y and the corresponding sparse solution \mathbf{c}_S to $\|X_{-\mathbf{y}}\mathbf{c}_S - \mathbf{y}\|_2 = 0$. Let $\mathcal{H}(\mathbf{y}, Q)$ be the plane through \mathbf{y} with orthonormal basis Q . Assume that there is a point $\mathbf{x}_n \in X_{T_y \setminus S_y}$ with $\angle(\mathbf{y}, \mathbf{x}_n) < \angle(\mathbf{y}, \mathbf{x}_{y,1})$ on the same side of $\mathcal{H}(\mathbf{y}, Q)$ as $\mathbf{x}_{y,1}$. Let S'_y be the set of indices of the points $\mathbf{x}_n, \mathbf{x}_{y,2}, \dots, \mathbf{x}_{y,d}$ and $\mathbf{c}_{S'}$ be the solution to $\|X_{-\mathbf{y}}\mathbf{c}_{S'} - \mathbf{y}\|_2 = 0$.

Proposition 2. *The solution $\mathbf{c}_{S'}$ of $\|X_{-\mathbf{y}}\mathbf{c}_{S'} - \mathbf{y}\|_2 = 0$ satisfies*

$$\|\mathbf{c}_{S'}\|_1 < \|\mathbf{c}_S\|_1. \quad (5)$$

Proof. The proof follows from Prop. 1. \square

Corollary 1. *The d points $\mathbf{x}_{y,1}, \dots, \mathbf{x}_{y,d} \in X_{T_y}$ closest to \mathbf{y} minimize $\|\mathbf{c}\|_1$ such that $\|X_{-\mathbf{y}}\mathbf{c} - \mathbf{y}\|_2 = 0$.*

The implication is as follows: The usual assumption is that the data uniformly distributed. In the context of sparse subspace clustering, this implies that data on the same subspace need be uniformly distributed (cf. (Soltanolkotabi and Candes 2012)). If this condition is violated, for instance because points fall into two well separated clusters, the support of the points of one cluster will not include points of the other.

This does not change the results in (Soltanolkotabi and Candes 2012) since the central Theorem 2.5 there rests upon the assumption that points on the same subspace are more or less evenly distributed. The idea in this work considers a particular violation of this assumption.

Robust l_1 -Norm: $\lambda > 0$

Define λ_1 and λ_2 such that $\lambda_1 + \lambda_2 = \lambda$. The idea here is to define an auxiliary line parallel to $\text{span}[\mathbf{x}_{y,1}]$ with distance λ_1 into direction of \mathbf{y} , and an auxiliary subspace parallel to $\text{span}[\mathbf{x}_{y,2} \dots]$ with distance λ_2 also into direction of \mathbf{y} . The geometry of this configuration is shown in the middle of Fig. 2.

Proposition 3. *The l_1 -norm of \mathbf{c} increases with the angle $\angle(\mathbf{y}, g(\mathbf{x}_{y,1}, \lambda_1))$ between \mathbf{y} and $g(\mathbf{x}_{y,1}, \lambda_1)$*

$$\|\mathbf{c}\|_1 \geq h_2(\eta, \lambda_2) \cdot \cot\left(f(\angle(\mathbf{y}, g(\mathbf{x}_{y,1}, \lambda_1))) + h_1(\rho, \lambda_2) - \frac{\pi}{2}\right). \quad (6)$$

Proof. The idea for the robust l_1 -norm estimator is to create auxiliary variables λ_1 and λ_2 so that $\lambda = \lambda_1 + \lambda_2$. These two are used to define two affine spaces parallel to $\text{span}(\mathbf{x}_{y,1})$ and $\text{span}(Q)$ with distances λ_1 and λ_2 such that the distances to \mathbf{y} are reduced by λ_1 and λ_2 . The function g constructs the affine space parallel to $\text{span}(\mathbf{x}_{y,1})$. The functions h_1 and h_2 modify the projection of \mathbf{y} not onto $\text{span}(Q)$ but its parallel affine space, and correct ρ , respectively. \square




Figure 2: Geometry of different l_1 -norm estimators. Left: noiseless case (Eq. (2) with $\lambda = 0$); middle: noisy l_1 -norm (Eq. (2) with $\lambda > 0$); right: lasso.

Lasso

The configuration for the lasso estimator

$$\min \|c\|_1 + \lambda \cdot \|X_{-y}c - y\|_2 \quad (7)$$

is very similar to the one shown for the robust l_1 -norm and is shown in the right plot of Fig. 2. We can now define a variable $\epsilon = \epsilon_1 + \epsilon_2$ that it absorbs the l_2 -error. Here ϵ_1 indicates the distance of an auxiliary line parallel to $\text{span}[y]$ with distance ϵ_1 and an auxiliary subspace parallel to $\text{span}[x_{y,2} \dots]$ with distance ϵ_2 . The difference to the robust l_1 -norm is that ϵ is now variable.

Due to the close similarity between the two estimators, Proposition 3 can easily be adapted to the lasso estimator:

Proposition 4. *The l_1 -norm of c is increases with the angle $\angle(y, g(x_{y,1}, \lambda_1))$ between y and $g(x_{y,1}, \lambda_1)$*

$$\|c\|_1 \geq h_2(\eta, \epsilon_2) \cdot \cot\left(f(\angle(y, g(x_{y,1}, \epsilon_1))) + h_1(\rho, \epsilon_2) - \frac{\pi}{2}\right). \quad (8)$$

Proof. The proof of Proposition 4 is identical as the one of Prop. 3 except that the admissible error $\epsilon = \epsilon_1 + \epsilon_2$ is now variable. \square

6 Selective Pursuit

Selective Dantzig Selector

Inspired by the Dantzig selector (Candes and Tao 2007; Qu and Xu 2015), $X^* = X^\top$, we notice that the product X^*X measures correlations to the subspace if we define a modified Dantzig selector

$$X^* = X_{S_y}^\top. \quad (9)$$

The matrix X_{S_y} consists of those columns of X that are the support of y . Since points on the same subspace have smaller distances to it than points from different subspaces, we can expect that the products X^*x_j , $j = 1, \dots, N \setminus \{\mathcal{I}(y), \text{supp}(y)\}$ to be large if $j \in T_y$.

The idea proposed here is as follows: The support set S_y of y usually consists of points from the set of neighbors of

y . Rather than stopping at this point, it is possible to select additional points, if their coherences with $\text{span}(X_{S_y})$ are large.

Let S_y^e be the extended support set of y which consists of the original support set, and the additionally selected points. If the modified Dantzig selector is now taken to be $X^* = X_{S_y^e}^\top$, the selection process becomes more and more influenced by noise. Because noise causes spurious singular values in X^* , we reduce their effect by scaling with

$$\rho = \text{trace}(X^*)^\top X^*. \quad (10)$$

Since the trace of the square of a matrix equals the sum of its singular values, this reduces the effect of the spurious singular values. Finally, new extended support vectors are chosen if for a threshold δ

$$\arg \max_j \frac{\|X^*x_j\|_2^2}{\rho} > \delta, \quad j \in \{1, \dots, N\} \setminus \{\mathcal{I}(y), S_y^e\}. \quad (11)$$

Subspace Selector

To further reduce the effect of spurious singular values, it is possible to fit an $d^{(l)} \approx \text{length}(S_y)$ dimensional orthonormal basis B to X^*

$$B = \arg \max \sum \|BX_j^*\|_2^2, \quad j = 1, \dots, \text{length}(S_y^e). \quad (12)$$

We can now choose additional points based on the Euclidean distance to the subspace $\text{span}(B)$ by

$$\arg \max_j \|P_B x_j\|_2^2 > \delta, \quad j \in \{1, \dots, N\} \setminus \{\mathcal{I}(y), S_y^e\}. \quad (13)$$

for a threshold δ . Here, P_B is the orthogonal projector onto $\text{span}(B)$.

7 Experiments

To experimentally demonstrate the effect of disconnectivity between clusters of points on the same subspace, we created two normally distributed point clouds on the unit sphere in




Figure 3: Relative connectivities between the two clusters. The x -axes correspond to angles between the clusters (cf. the explanation in Sec. 7), the y -axes to the relative connectivity ξ defined by Eq. (15). Left: no noise; middle: noise with $\sigma = 1\%$; right: noise with $\sigma = 2\%$. The dotted green line corresponds to the lasso-method, the dash-dotted red line to the variational l_0 -norm optimization (Logsdon et al. 2012), the dashed blue line to the proposed Selective Dantzig Selector, and the solid black line to the proposed Subspace Selector.

\mathbb{R}^3 with identical mean and standard deviation. Each clusters consists of 20 points. One was then rotated in steps of 5° from 0° to 180° . Afterwards, both point sets were multiplied by the same orthonormal 20×3 basis matrix. Each point was then normalized to length 1.

For each angle between the clusters, normally distributed noise was added to the data. We used standard deviations of $\sigma = \{0, 0.02, 0.03\}$ which amounts to 0%, 2%, and 3% noise. For each combination and noise magnitude, 10 trials were performed, i.e. the data was perturbed 10 times with different random noise.

To measure the connectivity between the two point sets, we analyzed the affinity matrix $A = |C| + |C^\top|$

$$A = \begin{bmatrix} A_{11} & A_{12} \\ A_{21} & A_{22} \end{bmatrix}. \quad (14)$$

The coefficients in the A_{11} and A_{22} blocks indicate the connectivity within each cluster whereas the coefficients in $A_{21} = A_{12}^\top$ indicate the connectivity between the different clusters.

We can therefore measure the connectivity ξ between the clusters by

$$\xi = \frac{1}{N} \sum_{i=1}^N \frac{\sum_{j=1}^{N_2} A_{12}(i, j)}{\sum_{j=1}^N A(i, j)} \quad (15)$$

The two algorithms proposed in Sec. 6 are compared against a lasso-estimator

$$\|\mathbf{c}_c\|_1 + \lambda \|Y_{-j} \mathbf{c}_c - \mathbf{y}\|_2 \quad (16)$$

which is optimized using CVX (CVX). We used only Eq. (16) because the results of Eq. (2) are almost identical. Further, an algorithm to estimate the l_0 -norm using a variational approach (Logsdon et al. 2012) was included into the comparison. This is motivated by the idea that the l_0 -norm is not

supposed to be effected by the bias that the l_1 -norm estimators have.

The results are shown in the plots in Fig. 3. The left plot corresponds to no noise, the middle one to a medium noise ($\sigma = 1\%$), and the right plot to strong noise ($\sigma = 2\%$). The x -axes correspond to angles between the clusters (cf. the explanation in Sec. 7), the y -axes to the relative connectivity ξ defined by Eq. (15). The dotted green line indicates the lasso-method (16), the dash-dotted red line the l_0 -norm optimization (16), the dashed blue line the proposed Selective Dantzig Selector, and the solid black line the proposed Subspace Selector.

As can be seen, the two proposed algorithms do not suffer as much from the gap between the two clusters as the other algorithms. Very surprisingly, the l_0 -norm estimator is also strongly affected.

Average affinity matrices computed by three different methods are shown in Fig. 4. Apparently, for the lasso estimator the two clusters cause almost disconnected subgraphs at an angle of 45° (left column in Fig. 4).

8 Summary and Conclusions

The topic of this work is sparse subspace clustering. The principle idea behind this class of algorithms is to approximate each data point by a linear combination of as few other data points as possible. This problem is often approached by solving l_1 -norm problems. The sparse coefficients are used to define edge weights of a graph. Computing a minimum cut through this graphs reveals which points are located on the same subspace.

Since there are few edges per vertex, it could be possible that vertices of a graph form disconnected subgraph although the corresponding points lie on the same subspace. In a previous work, the possibility for this to happen was answered negatively if the subspace dimension is 2 or 3. Four 4-dimensional subspaces, an example was given when dis-




Figure 4: Average affinity matrices $A = |C| + |C^\top|$ of the ten trials with their entries normalized to $[0, 1]$. The left column are the average affinities if the angle between the clusters is 0° . The middle column corresponds to angles of 10° , and the right one to 45° . (a) average affinities computed by the lasso estimator (16). (b) affinities compute by the variational estimator of the l_0 -norm (Logsdon et al. 2012). (c) affinities due to the proposed Selective Dantzig Selector.

connectivity occurs.

This work investigates the relaxed definition of not exact disconnectivity but *relative* disconnectivity. This means that all edge weights between two subgraphs are so low that the minimum cut either separates the two subgraphs, or – if the number of clusters is fixed – estimates a very erroneous result.

It is shown that this problem is caused by a gap between the points corresponding to the vertices of the subgraphs. In other words, if points on the same subspace fall into to separate clusters, then the two subgraphs necessarily have a so *low* connectivity that a successful clustering is no longer possible.

This problem does not invalidate the result published in (Soltanolkotabi and Candes 2012), since the authors explicitly state that their Theorem 2.5 requires well distributed points. The idea in this work is to consider the case that this assumption is not given.

Two different algorithms are proposed which are not as susceptible to the shown effect. Surprisingly, even an es-

timator of the l_0 -norm suffers from biased distributions of points.

References

- Candes, E., and Tao, T. 2007. The dantzig selector: Statistical estimation when p is much larger than n . *The Annals of Statistics* 35(6):2392–2404.
- CVX: *Software for Disciplined Convex Programming*. www.cvx.com.
- Elhamifar, E., and Vidal, R. 2013. Sparse subspace clustering: Algorithm, theory, and applications. *IEEE Transactions on Pattern Analysis and Machine Intelligence (TPAMI)* 35(11):2765–2781.
- Ho, J.; Yang, M.; Lim, J.; Lee, K.; and Kriegman, D. J. 2003. Clustering appearances of objects under varying illumination conditions. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, 11–18.
- Hong, W.; Wright, J.; Huang, K.; and Ma, Y. 2005. A multi-

scale hybrid linear model for lossy image representation. In *International Conference on Computer Vision (ICCV)*, 764–771.

Kanatani, K. 2001. Motion segmentation by subspace separation and model selection. In *International Conference on Computer Vision (ICCV)*, 586–591.

Li, C., and Vidal, R. 2015. Structured sparse subspace clustering: A unified optimization framework. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 277–286. IEEE.

Liu, G.; Lin, Z.; and Yu, Y. 2010. Robust subspace segmentation by low-rank representation. In *27th International Conference on Machine Learning (ICML)*, 663–670.

Logsdon, B.; Carty, C.; Reiner, A.; Dai, J.; and Kooperberg, C. 2012. A novel variational bayes multiple locus z-statistic for genome-wide association studies with bayesian model averaging. *Bioinformatics* 28(13):1738–1744.

Nasihatkon, B., and Hartley, R. 2011. Graph connectivity in sparse subspace clustering. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2137–2144.

Patel, V. M.; Nguyen, H. V.; and Vidal, R. 2013. Latent space sparse subspace clustering. In *IEEE International Conference on Computer Vision (ICCV)*, 225–232.

Qu, C., and Xu, H. 2015. Subspace clustering with irrelevant features via robust dantzig selector. In *Advances in Neural Information Processing Systems (NIPS)*, 757–765.

Soltanolkotabi, M., and Candes, E. 2012. A geometric analysis of subspace clustering with outliers. *The Annals of Statistics* 2195–2238.

Vidal, R.; Ma, Y.; and Sastry, S. 2005. Generalized principal component analysis (GPCA). *IEEE Transactions on Pattern Analysis and Machine Intelligence* 27(12):1945–1959.

Zhang, T.; Szlam, A.; Wang, Y.; and Lerman, G. 2012. Hybrid linear modeling via local best-fit flats. *International Journal on Computer Vision* 100(3):217–240.